A Comparative Study of Iterative and Conjugate Gradient Methods for Phase Retrieval from Fresnel Images

R Vincent, Physics Department, University of Bristol, UK

Fresnel Diffraction

The efficiency of practical methods for phase retrieval from the intensity distribution in Fresnel (near- field) images is investigated in the paraxial limit. The aim was to develop a robust and efficient algorithm for phase retrieval from aperiodic objects in the limit where phases may vary continuously by several multiples of 2ð across the sampled field. The equations for Fresnel propagation of a coherent wave from the object plane are well known, where the complex amplitude and phase distribution in the *x-y* plane at *z* is given by f_z. Then we have

$$f_z = f \otimes h \qquad \text{and} \qquad F_z = F H . \tag{1}$$

For continuous functions, the two versions of Eq.(1) are equivalent, being either the convolution of *f* with the Fresnel kernel *h*, or wave propagation of *F* (capitals represent the Fourier transform of a function). We have

$$h_{x,y} = \frac{\mathrm{i}}{\lambda z} \exp\left(\frac{-\mathrm{i}\pi\delta^2 \left((x-x_0)^2 + (y-y_0)^2 \right)}{\lambda z} \right) \tag{2}$$

$$H_{u,v} = \exp\left(\frac{+\mathrm{i}\pi\lambda z \left((u-u_0)^2 + (v-v_0)^2 \right)}{n^2\delta^2} \right)$$

Eqs. (2) are written for sampled arrays, where *ä* is the sampling interval, *ë* is the wavelength, *n* is the linear dimension of a square array, x_0 and y_0 are the image offsets in pixel units, with the equivalent beam tilt offsets given by u_0 and v_0. Hereafter, the defocus *z* is quoted in units of $ä^2/ë$. To conserve flux, the wave propagation version of Eq. (2) was used here. As discussed by Mas et al. [1] (see Fig. 1), the Nyquist sampling limit for $H_{u,v}$ with *z* in the numerator of the exponent implies that we must remain within the near-field limit, $|z| \leq n$.

The error was defined as the sum of squared intensity differences between calculated and observed Fresnel images. For display purposes, the normalised square root of the error function is plotted here. As discussed below, this definition may have significant consequences for the convergence of phase retrieval algorithms.

Iterative algorithm

The standard method described by Fienup [2,3] was used for phase retrieval, based upon the input-output algorithm (eqs. (22) and (25) in [2]) with $\hat{a} = 0.7$. Trials with other versions of the algorithm and different values for the feedback parameter showed no significant improvement. To combine data from several Fresnel images, the new estimates for f_0 were combined by taking the arithmetic mean before applying the feedback parameter to find the next input to the algorithm.

Conjugate gradient algorithm

It is essential to use analytic expressions for the partial differentials of the error functional with respect to the free parameters. These expressions were used in a standard conjugate gradient routine that converges towards a local minimum of the error function. Here the relevant equations are summarised; a more detailed description is given in [4]. For a real variable $å$, the partial differentials of $å$ with respect to f are written as a complex array. In the notation of Thiébaut and Conan [5] an array element with indices x,y is defined as

$$\left(d_f^{\,e}\right)_{x,y} = \frac{\partial e}{\partial\left(\mathrm{Re}\left(f_{x,y}\right)\right)} + \mathrm{i}\,\frac{\partial e}{\partial\left(\mathrm{Im}\left(f_{x,y}\right)\right)} \tag{3}$$

The partial differentials with respect to the transformed array F are written as

$$\left(D_F^{\,e}\right)_{u,v} = \frac{\partial e}{\partial\left(\mathrm{Re}\left(F_{u,v}\right)\right)} + \mathrm{i}\,\frac{\partial e}{\partial\left(\mathrm{Im}\left(F_{u,v}\right)\right)}\ . \tag{4}$$

It may be shown that the two arrays of differentials are a Fourier transform pair.
The definition for the error used here gives

$$d^{\,e}_{f\otimes h} = 4 f \otimes h\left(\left|f \otimes h\right|^2 - g_z^{\,2}\right) \tag{5}$$

where $g_z^{\,2}$ represents the observed intensities in the Fresnel image. To find the gradients of $å$ with respect to f, we take the Fourier transform of eq. (5), multiply by H^*, and then take the inverse transform. This procedure uses the identity $D^e_F = D^e_{FH}\, H^*$. The gradients of $å$ relative to the phases $ö$ are given by

$$d^e_j = \mathrm{Im}\left(d^e_f\ f^*\right) \qquad \text{where} \qquad f = \left|f\right| \exp\left(\mathrm{i}\, j\right) \tag{6}$$

The Fourier transform of d^e_j, written as D^e_F, gives the gradients of $å$ relative to the Fourier components of the phase array. Similar equations were used to find the error gradients with respect to z, u_0 and v_0 (see [4]).

Comparison of wave propagation and convolution

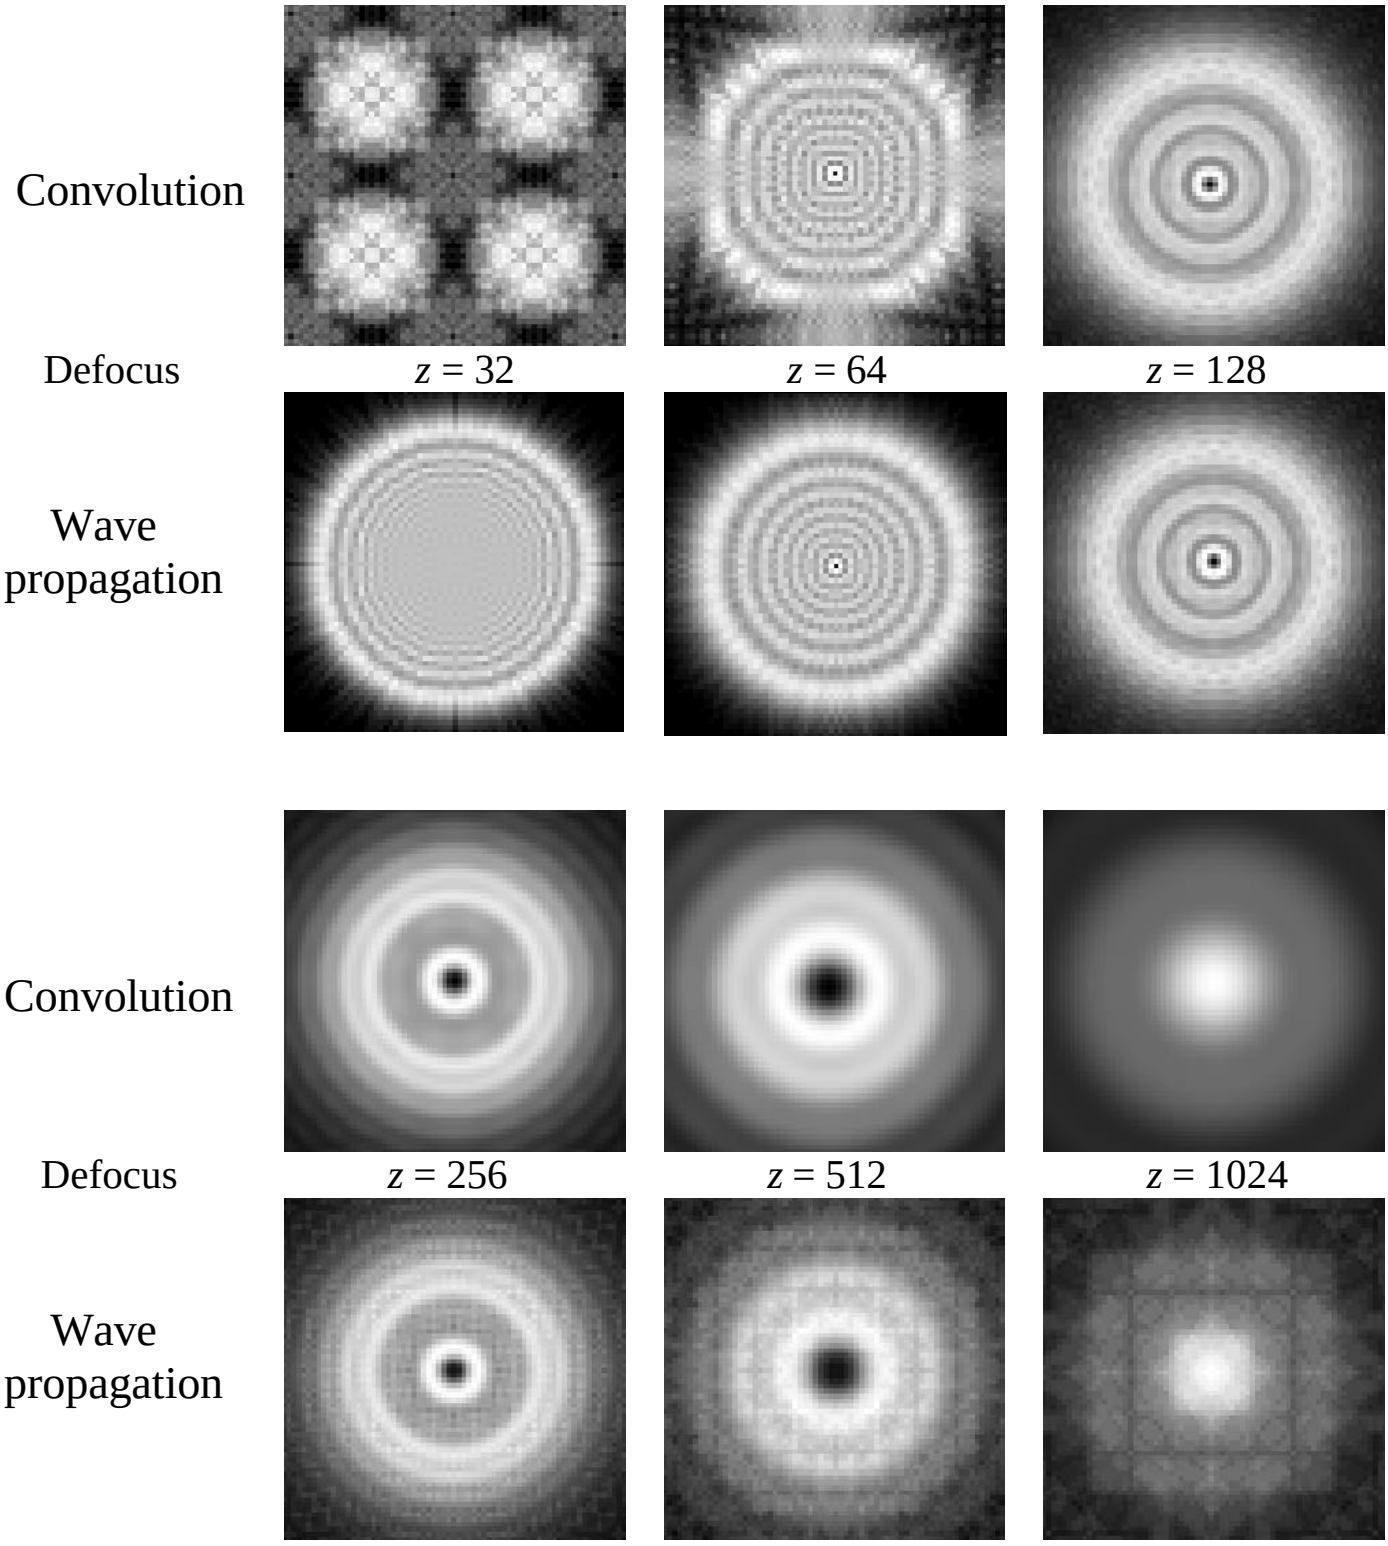

Fig. 1. Comparison of convolution and wave propagation methods for calculation of Fresnel images using FFTs. Array dimensions are n x n for a circular aperture radius $n/4$ (n =128). Defoci are given in units of $\ddot{a}^2/\ddot{e}$. As discussed by Mas et al., the Nyquist limit applied to the exponent in the Fresnel equations implies that calculations based upon wave propagation remain accurate only for $|z|$ n. Conversely, convolution is accurate only for $|z|$ n.

Leone and co-authors [6, 7] have shown that some considerable insight may be gained into the occurrence and avoidance of local minima if the error is represented as the sum of squared intensity differences. A path in f is defined by $f = f_0 + \ddot{e}\, n$, where f_0 is the (unknown) global minimum for $\mathring{a}$, n is another arbitrary vector and $\grave{\imath}$ is a scalar variable. The path remains linear if transformed into the space of f_z, but becomes quadratic in a space of the squared amplitudes for any z. A line search direction in f is parabolic in $|f_z|^2$. If the error function is defined as

$$e = \left\| |f|^2 - g^2 \right\|^2 + \left\| |f_z|^2 - g_z^2 \right\|^2, \tag{7}$$

then $\mathring{a}$ represents the squared distance between two points on a parabola. By substitution into eq. (7), an equation for $\mathring{a}$ is derived that is quartic in $\grave{\imath}$, where

$$\mathring{a} = \grave{\imath}^2 (a\,\grave{\imath}^2 + b\,\grave{\imath} + c). \tag{8}$$

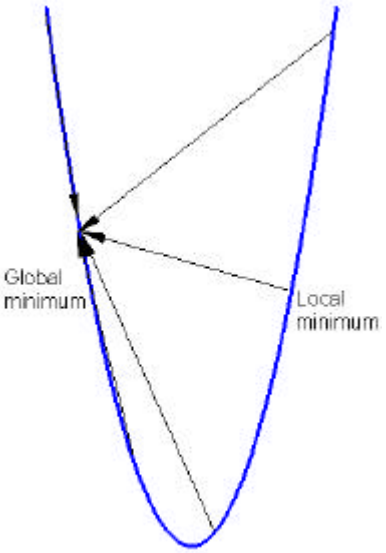

As shown opposite, the presence of a local minimum for $\mathring{a}$ depends upon the curvature of the parabola. The condition for a single minimum at $\mathring{a} = 0$ is given by

$$\frac{b^2}{ac} < \frac{32}{9} \tag{9}$$

Knowledge of a, b and c implies that the global minimum along the search direction in a conjugate gradient routine is found by solution of a cubic in $\grave{\imath}$. This produced a significant reduction in the processing time, but was not used in practice because the defocus and image offset parameters are non-linear in the space of f.

In eq. (9), both a and c are sums over squared terms, and therefore always positive definite. However, b is a sum over cross-product terms that are likely to be positive and negative. It follows that the best way to reduce the incidence of local minima is to increase the number of terms that contribute to a and c. This can be achieved by adding more images to the data representation space. The introduction of a compact support increases the size of terms beyond the support that contribute to c, reducing the probability of convergence to local minima that do not satisfy the support conditions. Convergence is improved by maximising the ratio of known to unknown parameters, achieved by using several images combined with compact support of the object. An alternative procedure is to impose compact support on f in inverse space, equivalent to bandwidth limit-ation. In practice, it was found that using the phases $\ddot{o}$ as free parameters produced very similar behaviour to the arguments outlined above. In particular, the introduction of progressive bandwidth relaxation for both iterative and conjugate gradient routines was often an effective method for promoting convergence when neither routine would converge under standard conditions.

Phase retrieval from random arrays

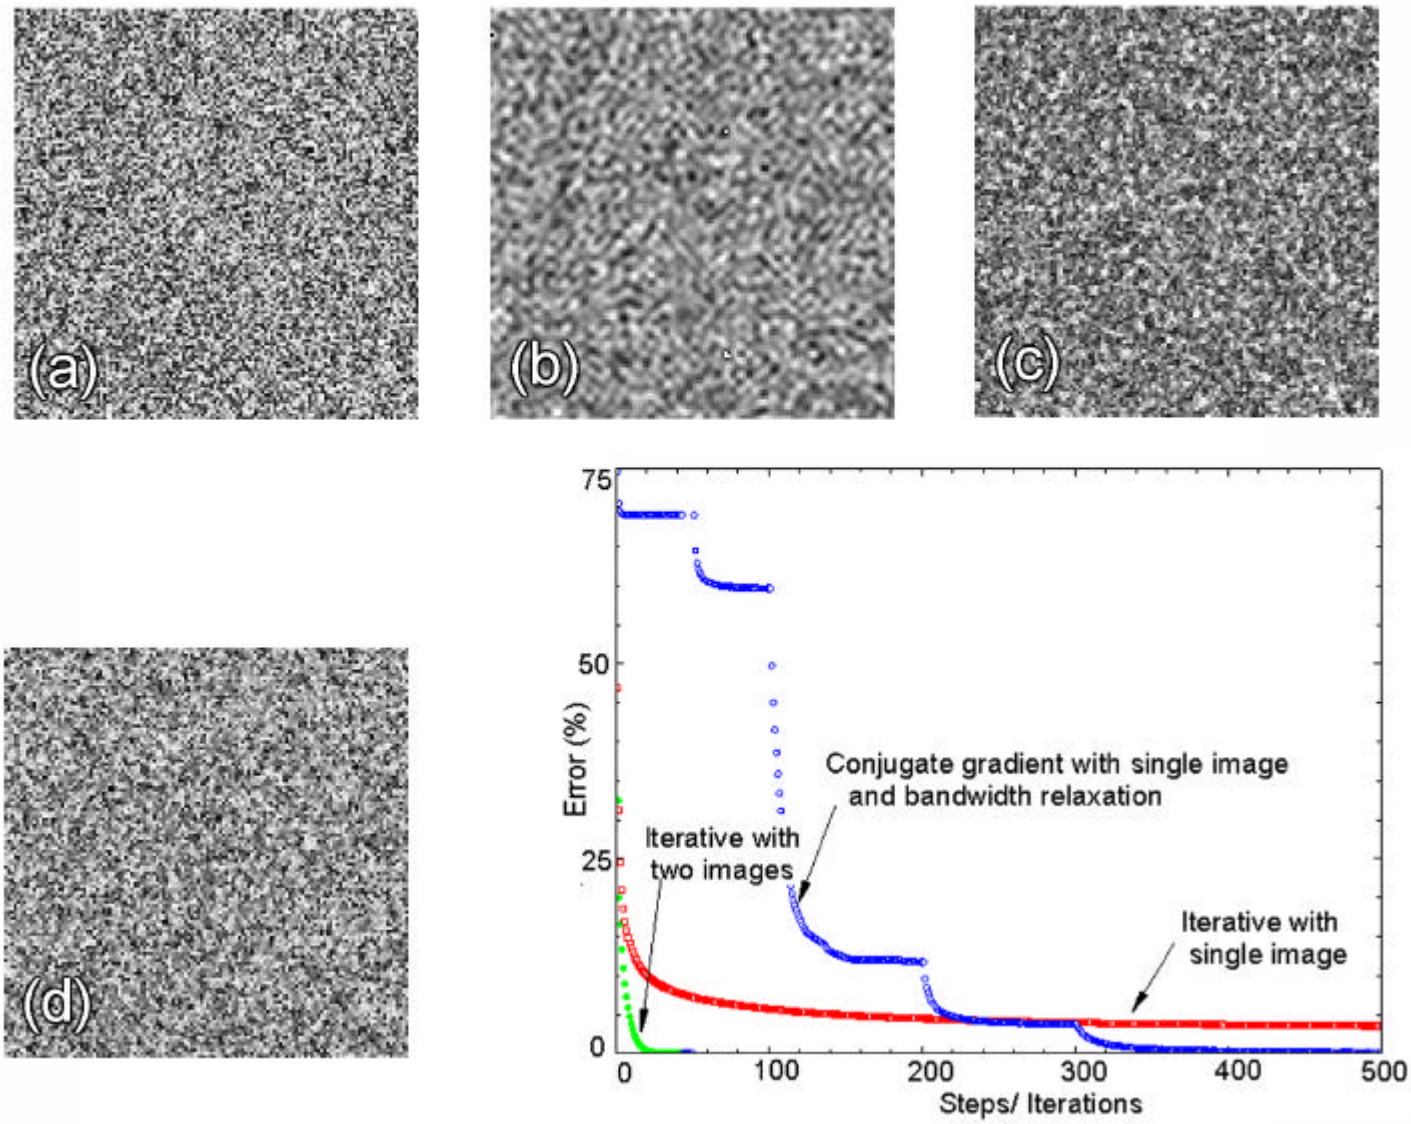

Fig. 2. The phases cannot be retrieved from an array with completely random amplitudes (0 to 1) and phases (0 to 2ð). However, reduction of the phase bandwidth by 50% along both axes (equivalent to a square aperture in Fourier space, or over-sampling of the image intensities) is sufficient to assure convergence from any random starting point.

Arrays (a) and (b) above show the object amplitudes and phases, respectively, where (c) is the Fresnel image at $z = 60$. The conjugate gradient routine converged exactly using only a single image and bandwidth relaxation. The equivalent iterative routine using a single image did not converge under any circumstances (see phase map (d) with 4% error). However, the use of two or more Fresnel images with the iterative routine was sufficient for rapid convergence if the starting phases were uniform. Convergence from a random array was very slow.

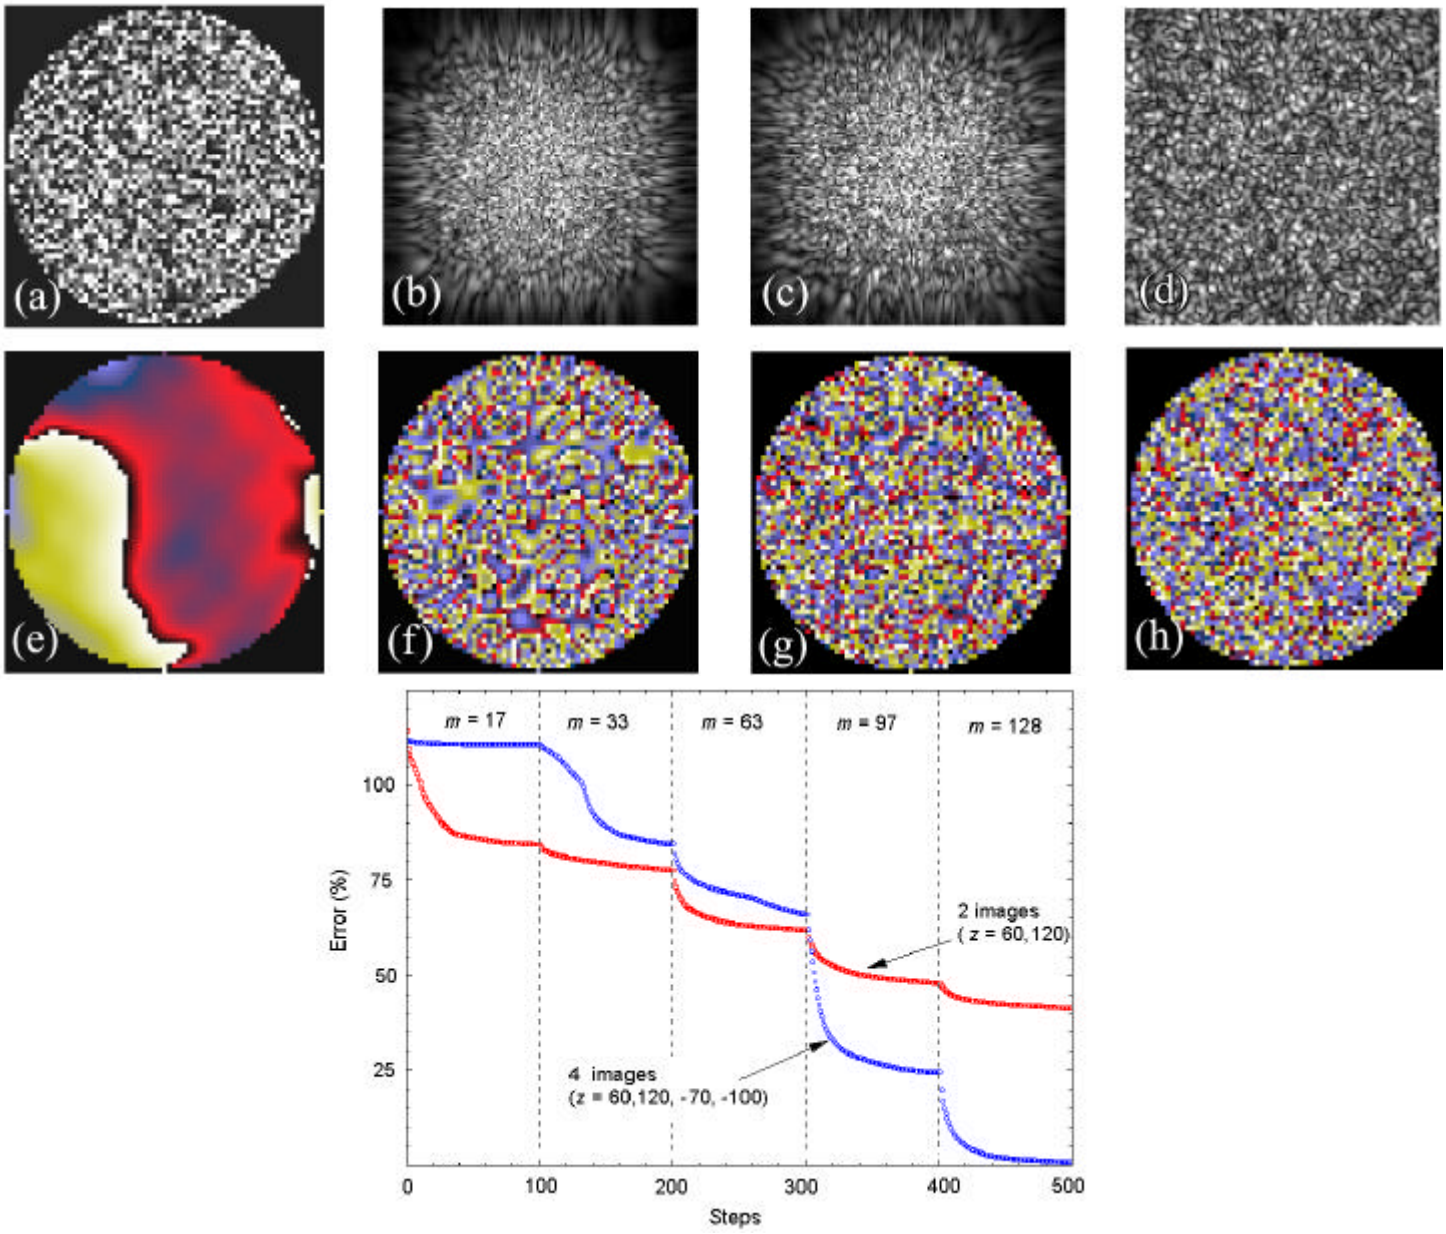

Fig. 3. Although random phases for an entire n^2 array cannot be retrieved from the Fresnel intensities (cf. Fig. 2), the use of a finite support reduces the number of free parameters and promotes convergence to the correct phase distribution. An example is shown here for a circular aperture, radius $n/4$ (n =128), where phases and amplitudes within the aperture are completely random (0 to 2δ and 0 to 1). The amplitude distributions in the object plane (65^2) and in three image planes (128^2) are shown in (a) to (d), where z = 0, 60, -70 and 120, respectively.

The iterative routine did not converge even after 5000 iterations using 4 images, with or without bandwidth relaxation. Typical errors were ~20%. The conjugate gradient routine did converge but only by using a combination of bandwidth relaxation and a minimum of 3 Fresnel images (see error curves above). The parameter m represents the reduced bandwidth, where the associated phase arrays for n = 17, 33, 67 and 128 are shown in (e) to (h), respectively. The requirement for several images at different defoci is equivalent to over-sampling of a single image, or bandwidth reduction as used for the object phase array in Fig. 2.

Retrieval of unwrapped phases

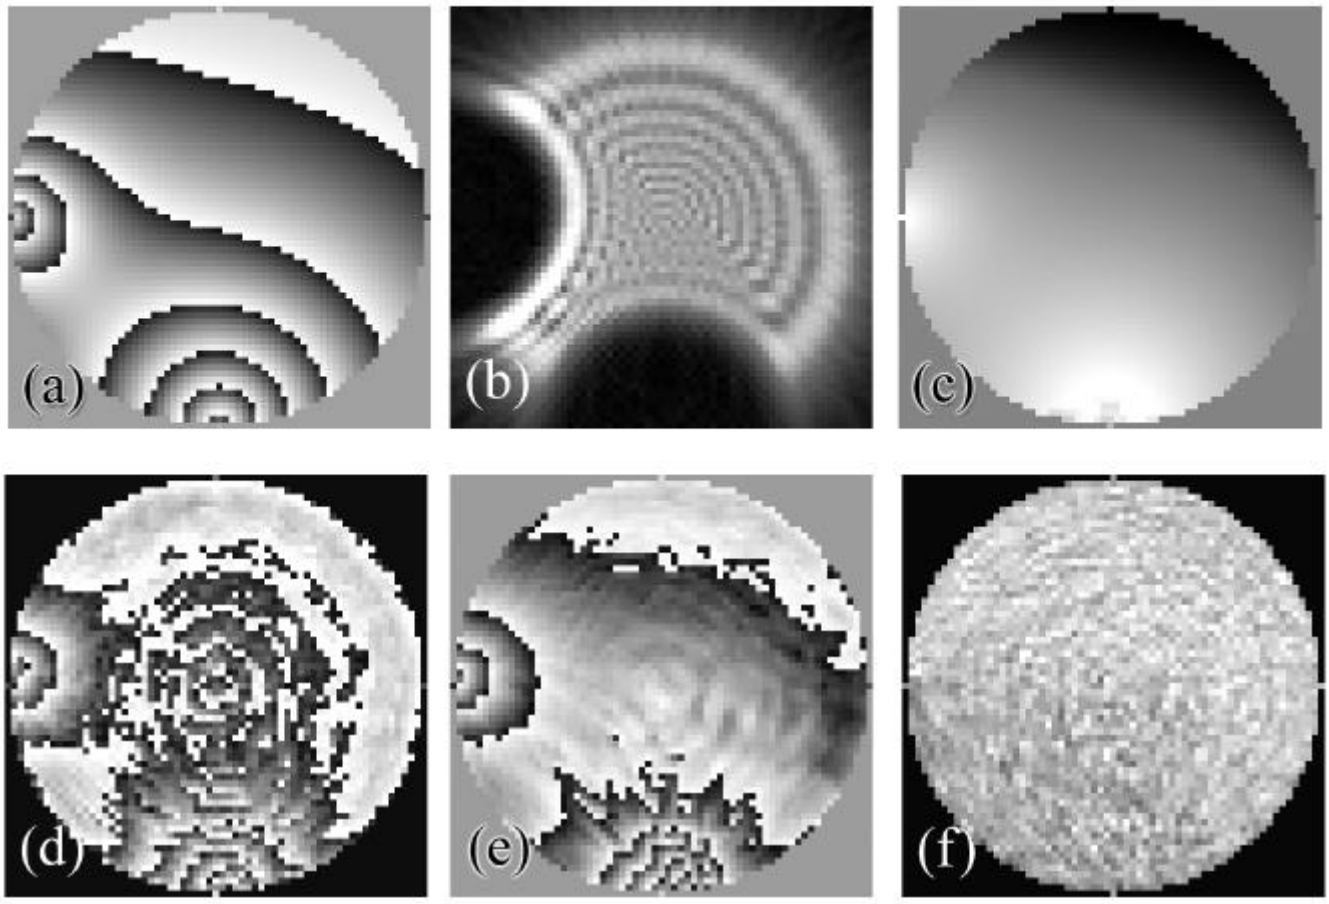

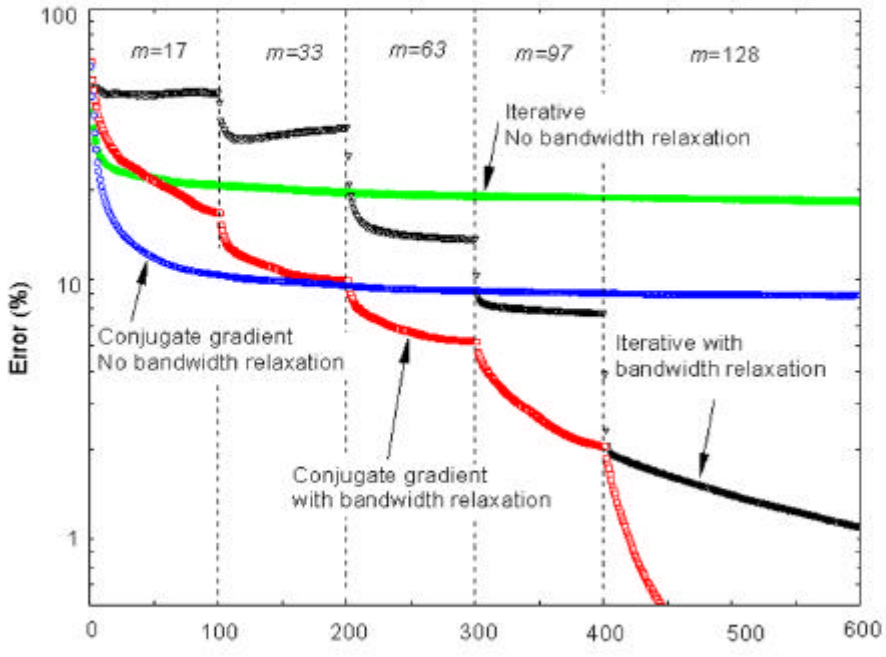

Fig. 4

Fig. 5. Linear traces through global and local minima (n=128)

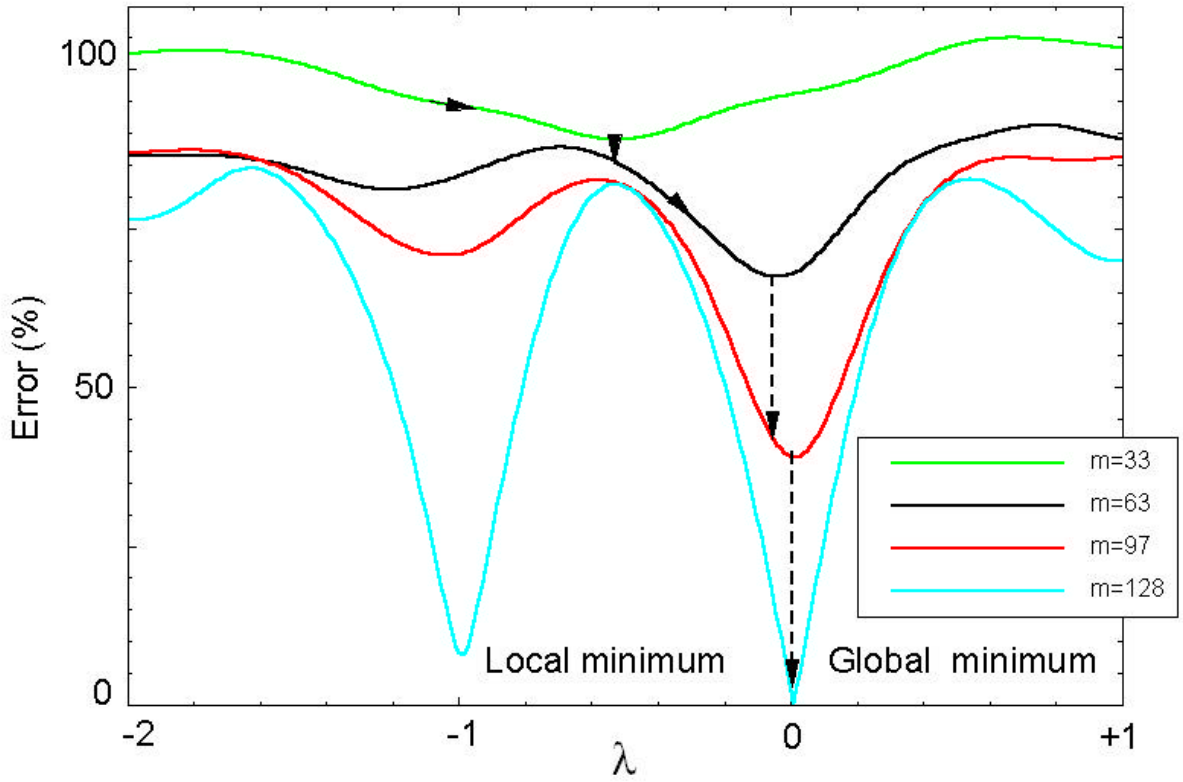

Figs. 4 and 5. For comparison with the random phase objects described in Figs. 2 and 3, the phase distribution in Fig. 4(a) was designed to resemble the phase shifts induced in an electron beam by electric or magnetic fields. The amplitude is constant within a circular mask, radius $n/4$ (n = 128). The unwrapped phase surface is continuous and varies smoothly over several multiples of 2ð. The single Fresnel image at z = 60 shown in Fig. 4(b) is used for phase retrieval

In the absence of bandwidth relaxation, neither algorithm converged to the correct phase distribution. The conjugate gradient routine was trapped in an apparent local minimum with an error of 8% (Fig. 4(d)). Likewise, the residual error for the iterative routine was 15% after 5000 iterations. The associated phase and amplitude maps are shown in Figs. 4(e) and (f).

However, both algorithms converged rapidly when progressive relaxation of the phase bandwidth was used. The reduced bandwidths (m = 17, 33, 63 and 97) are marked in Fig. 3. Significantly, the phase map recovered by the conjugate gradient algorithm was the unwrapped version (Fig. 4(c)), produced by continuous deformation of the phase surface.

Variation of the error along a line scan in phase space intersecting the global and local minima (Figs. 4(a) and (d), respectively) is shown in Fig. 5, together with equivalent scans for the phase maps with reduced bandwidths. Significantly, a reduction in the number of phase parameters eventually removes the local minimum. The dashed lines indicate the path followed by a gradient descent algorithm confined to the line scan direction.

Phase retrieval of text objects

Object intensity | Object phase | Image (z = 60) | Image (z = 120)

Start ($å$ = 46%) | m = 17 ($å$ = 27%) | m = 33 ($å$ = 19%) | m = 63 ($å$ = 12%)

m = 97 ($å$ = 7%) | End ($å$ = 0) | Single image (CG) | Single image (Iterative)

Fig. 6. The text object used here to test convergence is both easily recognised and includes large changes of amplitude and phase at the Nyquist sampling limit. The object amplitude and phase is shown in (a) and (b), together with two defocused images in (c) and (d). The amplitude is reduced by 50% within letters and the phase is shifted by $ð/2$. It was essential to use two images for rapid convergence of both iterative and conjugate gradient routines, although bandwidth relaxation was unnecessary if the initial phases were uniform. The progressive increase of resolution using bandwidth relaxation is shown in (e) to (j). If only a single image was used, convergence became extremely slow for both routines (see (k) and (l) after 500 and 5000 iterations, respectively). However, the residual errors for the phase maps shown in (k) and (l) were very small (typically 0.1%), demonstrating that a single parameter measure of the error is not always a reliable guide to image quality.

Phase shifts induced by a contaminated condenser aperture

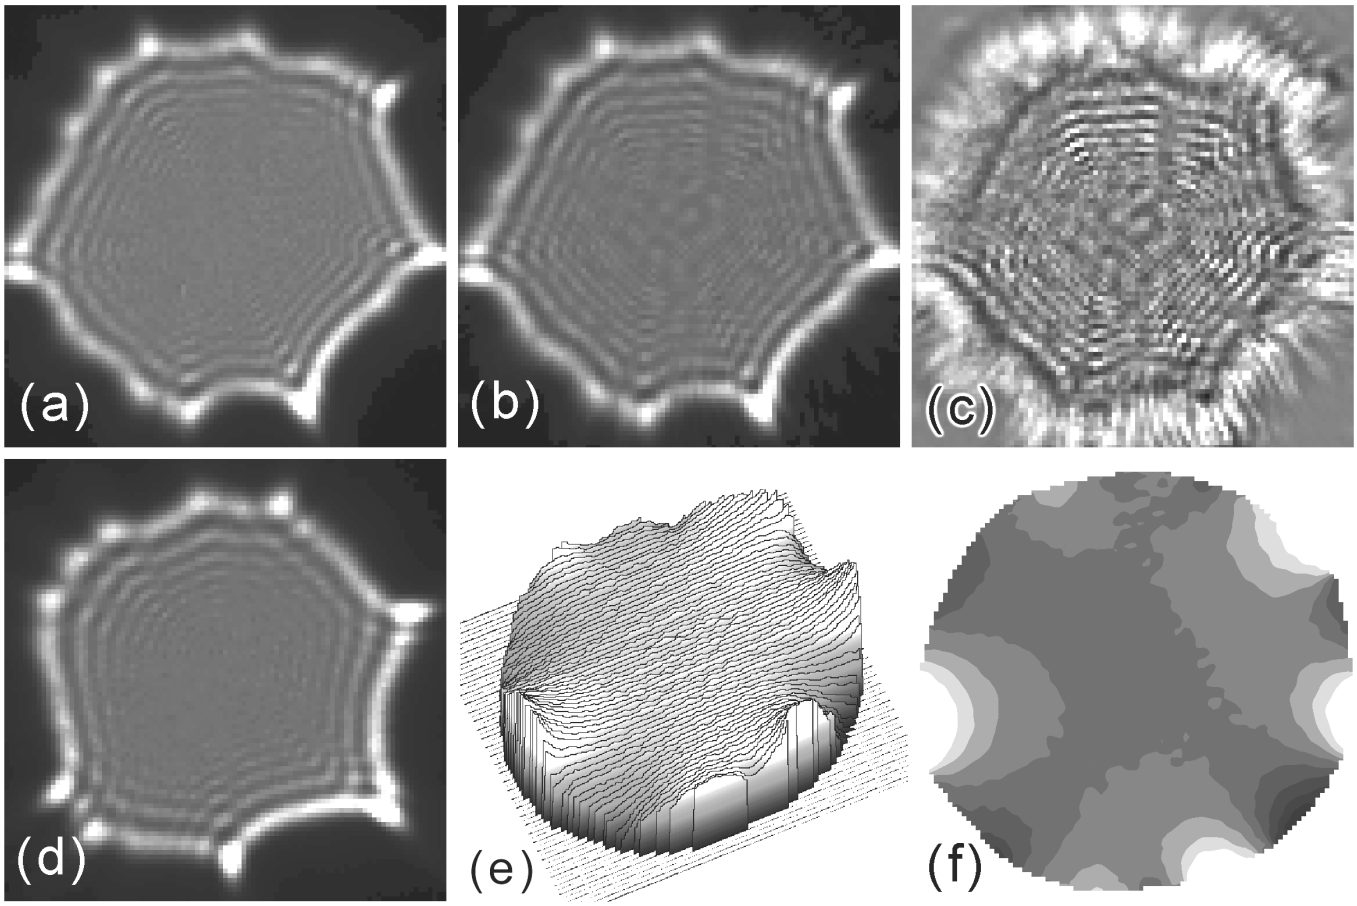

Fig. 6. Example of phase retrieval using defocused images of a contaminated 10 ìm condenser aperture illuminated by a coherent electron beam. Experimental under- and overfocused images of the aperture are shown in (a) and (d), respectively.

The images were corrected for small shifts in magnification and rotation relative to the object plane, and then phases were retrieved with the conjugate gradient routine using bandwidth relaxation. The final error for the two images was about 5%. For comparison with (a), the corresponding calculated image is shown in (b), together with the modulus of the difference map in (c). Differences between (a) and (b) are attributed to partial coherence in the incident beam, and consequent overfitting of the higher intensity fringes.

The phase map in the aperture plane is plotted in an isometric view (e) and as stepped contours at intervals of 1 rad in (f). The total phase shift is about 10 rad, equivalent to 1.5 wavelengths. The phase surface is almost flat near the centre of the aperture. Around the perimeter the phase is both advanced and retarded, attributed to negatively charged patches of contaminants balanced by positive charge in the aperture strip.

Discussion

Iterative algorithms are always considerably faster than any gradient descent algorithm that requires repeated calculations of the error function. As mentioned above, direct solution of a cubic equation for parabolic search directions in the data space did reduce the processing time by at least 50%. However, it was preferred to retain the flexibility of using the phases as free parameters together with the defoci and image offsets. It is worth noting that convergence is identical if the complete set of Fourier phase components are used as free parameters, being equivalent to a different set of orthonormal basis vectors for the same solution space. The use of progressive bandwidth relaxation was based upon the idea that a severely reduced bandwidth eliminates most of the local minima, and increases the probability of convergence to the global minimum for a restricted model of the phase surface. Naturally, the error is not zero because the model is incorrect. However, it is probable that the new starting point is within the catchment zone for the next global minimum when the bandwidth is increased. In this way, the search path is guided towards the correct minimum for an unrestricted bandwidth. The error curves in Fig. 5 illustrate this behaviour for a fixed line search direction.

Factors that may affect convergence include:

- Under-sampling of the image intensities.
- Extent and shape of the support.
- Number of images available (at suitable defocus intervals).
- Natural bandwidth of the object amplitudes and phases.
- Continuity of the unwrapped phase surface.
- Relevance of bandwidth relaxation.
- Starting point (random vs uniform phases).

Considering the simulated examples, it was not unexpected that an object with completely random amplitudes and phases and no support failed to converge under any circumstances. It seems probable this behaviour must be attributed to under-sampling of the images, given that a phase modulation with wavelength 2*ä* in the object is expected to generate intensity modulations with wavelengths close to *ä* in a near-field image. If *ä* is the sampling interval, the images are under-sampled. This was confirmed (Fig. 2) by halving the phase bandwidth. Both routines converged, although the iterative method required two images. The conjugate gradient routine required only a single image combined with bandwidth relaxation. Alternatively, the use of a circular support (radius *n*/4) with completely random phases and amplitudes was sufficient for the conjugate gradient routine to converge, although a minimum of 3 images and bandwidth relaxation were required (Fig. 6). The iterative routine did not converge.

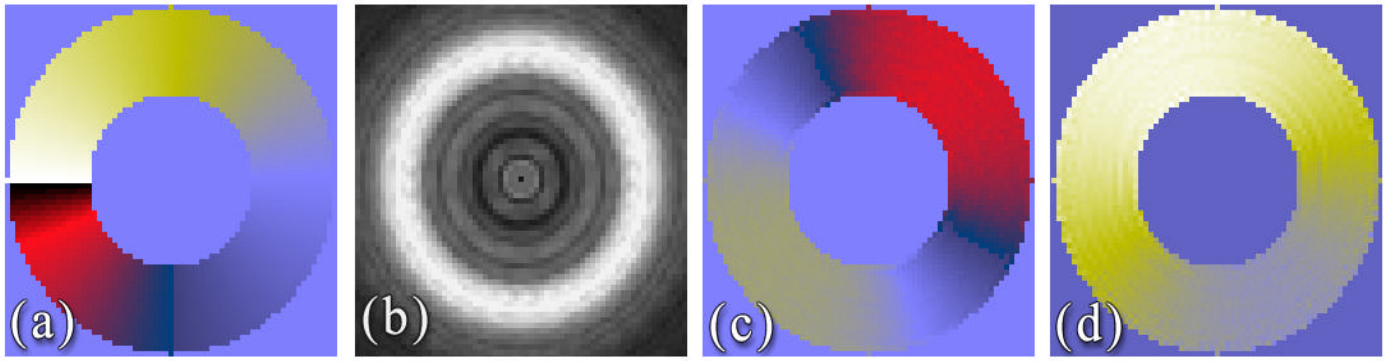

Fig. 7. (a) Object with annular support, constant amplitude and a helical phase surface increasing by 2ð for each revolution. As expected, all images show a dark spot at the centre (b). Both routines failed completely to converge under all circumstances. Examples of local minima are shown in (c) and (d).

The continuous unwrapped phase surface used in Fig. 4 was retrieved only when bandwidth relaxation was used. Otherwise, both algorithms became trapped in local minima. The text object (Fig. 6) contained short wavelengths, implying partial under-sampling of the image intensities, combined with large areas of constant phase and amplitude. For both routines, bandwidth relaxation was not strictly necessary, but it was essential to use two images to compensate for under-sampling. Convergence was very slow if a single image was used. It seems that the bandwidth relaxation procedure encourages convergence if an object contains a balanced mixture of wavelengths and is not under-sampled. It is especially useful if the phase surface is continuous and extends over several multiples of 2ð. Generally, there were no phase distributions that could be recovered only by the iterative routine, although there were several converse examples. The process of bandwidth relaxation encourages continuous deformation of an initially smooth phase surface. The example shown in Fig. 7 was designed as a severe test because the phase surface is both continuous and multiply valued. Neither routine converged.

References

[1] *Fast algorithms for free-space diffraction patterns calculation.*
D. Mas, J. Garcia, C. Ferreira, L. M. Bernardo and F. Marinho, Optics Comm., 1999, **164,** 233-245.

[2] *Reconstruction and synthesis applications of an iterative algorithm.*
J. R. Fienup, SPIE , 1981, **373,** 147-160.

[3] *Reconstruction of a complex-valued object from the modulus of its Fourier transform using a support constraint.* J. R. Fienup, J. Opt. Soc. Am. A, 1987, **4**, 118-123.

[4] *Phase retrieval in TEM using Fresnel images.*
R. Vincent, 2001, Ultramicroscopy, submitted for publication.

[5] *Strict a priori constraints for maximum-likelihood blind deconvolution.*
E. Thiébaut and J.-M. Conan, J. Opt. Soc. Am. A**,** 1995, **12**, 485-492.

[6] *Reconstruction of complex signals from intensities of Fourier transform pairs.*
G. Leone, R. Pierri and F. Soldovieri, J. Opt. Soc. Am. A**,** 1996, **13**, 1546-1556.

[7] *Radiation pattern evaluation from near-field intensities on planes.*
T. Isernia, G. Leone and R. Pierri, 1996, IEEE Trans. Antenna and Prop., 1996, **44**, 701-709.